

ABSTRACT

The present invention provides a method for copying data through a virtualized storage system using distributed table driven (I/O) mapping. In a system having a virtual disk (the "original disk"), a persistent mapping table for this virtual disk exists on a controller, and volatile copies of some or all entries in this mapping table are distributed to one or more more mapping agents. The method of the present invention creates a new virtual disk mapping table that has the exact same entries as the mapping table as the original virtual disk. The new snapshot disk then shares the same storage as the original disk, so it is space efficient. Furthermore, creating new snapshot disk involves only copying the contents of the mapping table, not moving data, so the creation is fast. In order to allow multiple virtual disks to share storage segments, writes to either the original virtual disk or the snapshot copy cannot be seen by the other. Therefore, in addition to simply copying the mapping table, both the original and snapshot disk mapping table must also cause writes to these disks to be handled specially. Finally, any changes to the original disk mapping table stored in the controller must be coordinated with the volatile copies stored in the mapping agent in such a way so that all hosts see a consistent view of that disk.